DTIC FILE COPY

(4)

AD-A214 286

Gaze Controls with Interactions and Delays

Christopher M. Brown

Technical Report 278
March 1989

UNIVERSITY OF
ROCHESTER
COMPUTER SCIENCE

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER 278 / 2. GOVT ACCESSION NO	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Gaze Controls with Interactions and Delays	5. TYPE OF REPORT & PERIOD COVERED Technical Report
	6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Christopher M. Brown	8. CONTRACT OR GRANT NUMBER(s) DACA76-85-C-0001
9. PERFORMING ORGANIZATION NAME AND ADDRESS Computer Science Department, 734 Computer Studies Bldg, University of Rochester. Rochester, NY 14627	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS D. Adv. Res. Proj. Agency, 1400 Wilson Blvd, Arlington, VA 22209	12. REPORT DATE March 1989
	13. NUMBER OF PAGES 19
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office) US Army ETL, Fort Belvoir, VA 22060	15. SECURITY CLASS. (of this report) Unclassified
	15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Distribution of this document is unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

None.

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

predictive control. active vision, gaze control

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

Five control systems loosely corresponding to primate saccadic, vergence, pursuit. vestibulo-ocular, and head control operate on a simulated two-eyed robot head maneuvered by a robot arm. The goal is to get some qualitative understanding of the interaction of such reflexes under various assumptions. The simulation is meant to be relevant to U. Rochester's robot. Thus it incorporates kinematics of the robot head but assumes a "tool-coordinate" system available to robot arm commands. so that arm kinematic calculations are unnecessary. Dynamics are not modeled, since they are handled by the commer-

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20. ABSTRACT (Continued)

cial controllers currently used in the Rochester robot. Even small delays render the effect of delay-free controllers unstable, but multi-delay version of a Smith predictor can cope with delays. If each controller a ts on the predicted system and ignores other controllers, the situation is improved but still potentially unstable if controllers with different delays act on the same control output. The system's performance is much improved if controllers consider the effect of other controllers, and the resulting system is stable in the presence of a certain amount of stochastic disturbance of control delays and inputs, and also in the presence of systematic error arising from inaccurate plant and world models.

GAZE CONTROLS WITH INTERACTIONS AND DELAYS

Christopher Brown
Computer Science Department
University of Rochester
Rochester, NY 14627

ABSTRACT

Five control systems loosely corresponding to primate saccadic, vergence, pursuit, vestibulo-ocular, and head control operate on a simulated two-eyed robot head maneuvered by a robot arm. The goal is to get some qualitative understanding of the interaction of such reflexes under various assumptions. The simulation is meant to be relevant to U. Rochester's robot. Thus it incorporates kinematics of the robot head but assumes a "tool-coordinate" system available to robot arm commands, so that arm kinematic calculations are unnecessary. Dynamics are not modeled, since they are handled by the commercial controllers currently used in the Rochester robot. Even small delays render the effect of delay-free controllers unstable, but multi-delay version of a Smith predictor can cope with delays. If each controller acts on the predicted system and ignores other controllers, the situation is improved but still potentially unstable if controllers with different delays act on the same control output. The system's performance is much improved if controllers consider the effect of other controllers, and the resulting system is stable in the presence of a certain amount of stochastic disturbance of control delays and inputs, and also in the presence of systematic error arising from inaccurate plant and world models.

INTRODUCTION

Behaving, actively intelligent (mechanical or biological) systems must manage their computational and physical resources in appropriate ways in order to survive and to accomplish tasks. At Rochester we are building an integrated actively intelligent system that incorporates abstract reasoning (planning), sensing, and acting [Bro88]. The *active intelligence* paradigm we shall exploit incorporates the following ideas.

1. A hierarchy of control, so that the highest cognitive levels can reason in terms of *what* they want done rather than *how* to do it in detail. This hierarchy should extend throughout the system.

2. At the lower levels, the control hierarchy ends with visual and motor skills or *reflexes*. These capabilities are cooperative but to some extent independently controllable. Some are always running, and they form the building blocks on which more complex behavior is [illegible]. Examples are tracking targets to minimize motion blur or redirecting gaze as a result of attentional shifts.

3. Part of the job of low-level visual capabilities is to present perceptual data, such as flow fields or depth maps, to higher-level visual processes. Low-level processes can often benefit from knowledge of self-initiated motion on the part of the sensing entity. They can often be built on the low-level control capabilities.

We currently have a nine degree of freedom robot body-head combination controlled by a Sun computer interfaced over a serial line to a VAL-II robot control system, and over a VME bus to the three eye motor controllers. The visual input is processed by a pipelined image processing system. The system has been used in several promising demonstrations of considerable complexity in depth-map creation and vergence

([BOS88,OPS89]). It has also been used for some simple but effective real-time applications in tracking and fixation.

What has been missing so far has been the cooperation of several modes of control, or the operation of several at once. In the work reported below, a simulation of the robot head and eyes is used to examine the effects of different styles of interaction between certain control capabilities that we have implemented (such as tracking) or anticipate using (such as using eye movements to compensate for head movements).

The simulation software is based on the actual robot head kinematics, and has provided a flexible tool for investigating the interaction of different control methods and different types of control interaction.

THE MODEL OF HEAD AND IMAGING

The simulator geometry can capture all the essentials of the Rochester robot [Bro88,BR88] (including the annoying "non-spherical" geometry of the camera pans and tilts). It allows geometric parameters to be changed to explore the effects on error and the possibility of adaptative control. The robot arm is not modeled: rather the model abstracts it to a single eye-support platform that can be postioned arbitrarily in space with six degrees of freedom: three in position, three in orientation. On the model head is a modelled tilt capability that affects both cameras, and each camera has a modelled pan capability. The geometry of the offsets of the various axes in these links are variable, and incorporate the geometrical complexity of the real system. The simulated mechanism is massless; this reflects the effective behavior of our current hardware system when viewed from its high-level control operations. The independent control of the camera pans allows us to model modern theories of saccadic and vergence systems; heads with mechanical vergence capability need one fewer motor but must use older models of these systems.

The camera models incorporate point projection with fixed focal length, as well as a "foveal-peripheral" distinction by which the location of imaged points is less certain, outside a small foveal region, depending on the off-axis angle of the target being imaged. The target itself is a single point in 3-D space, moving under dynamical laws. The experiments below were often carried out with the target point in orbit about an invisible "black hole" – thus the target followed an elliptical path. In other experiments the target moved in a straight line. In some of the experiments involving delays the target was stationary but the robot moved in X, Y, and Z, thus creating a perceived target motion, but one due to factors under robot control.

It is assumed that the imaging system knows the distance to the target (in real life, this distance may be derived from binocular stereo, apriori knowledge, any of a number of monocular distance cues, kinetic depth calculations, etc.). It is assumed that, for each eye, the instantaneous retinal velocity of the target is known (i.e. the vector difference between its position in the current image and its position in the last image). Other than that, the system only knows the left and right image (x,y) location of the target's image. Of course the target's image position and hence image velocity is perturbed by uncertainties arising from the blurriness of peripheral vision, should the target not be foveated. There is a further provision to add uniform noise to the target's imaged position –, this can model quantization noise, or be used to approximate process noise in the target's motion.

THE MODEL OF CONTROL

ZERO DELAY CONTROL

The input to the control systems is usually based on quantities that can be inferred from vision (e.g. the (x,y) position of the target, which should be driven to (0,0), or target disparity between the two eyes which should

be driven to 0). Some control inputs arise from the robot's "proprioception" (e.g. the amount the cameras are panned or tilted from their null position), and some is from other control signals (when one control is to null out the effects of another). The simulation has controllable output parameters corresponding to one set of VAL-II robot control parameters (the VAL-II "tool coordinate system") for the head: its X,Y,Z position and A,B,C orientation. Also there is direct control over the pans (independent for left and right) and tilt (common) of the two cameras. In every case the outputs of controls are velocity commands to the nine degrees of freedom in the system, reflecting one simple form of our current interface to the motor controllers.

The basic control loops that manage the system are loosely inspired by the primate visual system. However, most assumptions and technical decisions have been made either for the sake of simplicity or to mimic our robot rather than for the sake of faithfully modelling known biological systems or optimal mechanical systems (see the Discussion section below). Still, one of the major design goals is that the system can support more detailed control models. Most of the loops have several parameters, such as the proportional, integral, and derivative (PID) constants of their controllers, and their delays and latencies. Delay means the amount of time after a commanded motion before it commences — this is often called latency in the literature. Latency is how long it takes the command to complete: it is another time constant that indicates both how soon another command can be accepted, or how long the command will be affecting the controlled (velocity) variables. In all the work so far, only saccades have latency greater than unity. In the robot system the delay corrreponds to how long it takes the mechanical system to respond to a motion ordered from a high software level, and the latency reflects how long it takes to complete a command. The assumption is of control delay, not sensor delay: that is, we assume that "sensors" (visual or robot- and eye-control motor states read from their controllers) are available to the system immediately, without delay, and thus reflect the true state of the world. (Our analysis and the algorithms extend to the case that the sum of control and sensor delays is constant for any controller.)

There are five separate control systems.

1. Saccade: fast slewing of cameras to point in commanded direction. Saccades are modelled as open loop, though in primates there are "secondary" saccades that correct errors in initial saccades. The saccadic system tries to foveate the target and to match eye rotations to the target velocity so as to be tracking the target as soon as the saccade is completed. Current opinion is that the saccadic system is aware of the 3-D location of the target, not just the location of its retinal image. However, in the implementation used for the experiments below, saccades operate with retinal locations and velocities, not 3-D locations or distance. The left eye is dominant in the system. The saccade aims to center the target image on the fovea of the left eye; the right eye is panned by the same amount (and of course tilted by the same amount for mechanical reasons). Thus the saccade maintains the current vergence angle. It is implemented as a constant-speed slewing of all three pan and tilt axes, with one of them attaining a system constant maximum velocity. The slewing continues until the target should be foveated (it my not be due to peripheral blurring or other noise), at which time the system is left with eye velocities that match the perceived target motion before the saccade. The saccadic system is characterized by its maximum velocity and its delay.

2. Smooth Pursuit: tracking a moving target. This is a "continuous" activity as opposed to the discontinuous saccadic control activity. The error here is target position in the left eye, (which should be (0,0)), and the commands are pan and tilt velocities to the left eye. The pursuit system has delay, latency, and PID control. In both the saccadic and smooth pursuit systems modeled here, there is strict (exclusive) left-eye dominance.

3. Vergence: the vergence system measures horizontal disparity between the target position in the left and right eyes, and pans the right eye to reduce it. The vergence system has delay, latency, and PID control.

4. Vestibulo-Ocular System: the VOR system is open loop in the sense that its inputs come from the head positioning system and its outputs go to the eye positioning system. Its purpose is to stabilize

eyes against head motion, and its inputs are the control signals for head position (XYZ velocities, ABC angular velocities). It also uses the distance of the target, since that affects the appropriate response. The VOR should ideally be implemented by inverse kinematics, to which the current implementation (and presumably the neural one) is an approximation. Its output is commands to the pans and tilt controls to null out the apparent target motion caused by head motion. It is characterized by delay, latency, and open loop proportional gain.

5. Platform Compensation: This system is a head-control, not gaze-control system. These systems are known to interact in subtle and complex ways, but this particular reflex simply attempts to keep the eyes "centered in the head", so that the camera pans or tilts are kept within "comfortable" mechanical ranges. The "comfort function" is a nonlinear one $x/((x - xmax)^2)$, where x is the average pan angle (to control head "yaw" movements) or the tilt angle (to control head "pitch" movements). In either case $xmax$ is the mechanically imposed limit of the system. This reflex is open loop (eye position affects head position), with delay, latency, and open loop proportional gain

The system has the capability of operating in two modes: *smooth pursuit* and *saccade*. In smooth pursuit mode, the VOR, platform compensation, pursuit, and vergence systems are left running. In saccade mode, other controls may be disabled. This allows modelling the effects of turning off vergence, head compensation, tracking, *etc.* during saccades. Ultimately it seemed best only to turn off tracking during saccades, but other combinations are demonstrated below.

The delays and latencies are implemented with a command pipeline, in which the commanded changes in velocities are entered opposite the time in the future they are to take effect. Time is discretized to some level, called a *tick* henceforth. A larger delay results in entry of the corresponding command further in the future. Latencies are implemented by dividing the commanded change between as many discrete time periods as necessary to spread the effect over the latency. The pipeline thus is indexed by (future) time instant, and it has entries that hold the commanded velocities for the six head degrees of freedom and three camera degrees of freedom. Each instant also has an entry corresponding to its mode (saccadic or pursuit). The pipeline is implemented as a ring buffer.

For the delay-free case, the control architecture is strictly independent. That is, controllers are ignorant of each other's effects, and the combination of control effects is modeled by all controllers incrementing or decrementing a common control register (indicating some motor velocity setting). All increments and decrements are made to the current value that is there already, which perhaps is nonzero because of input from another reflex. Thus the control commands are summed in the simplest possible way, as if each control system's output were a D.C. voltage and all the outputs were soldered together at the effector motor's input.

The saccadic system shuts down the pursuit system in the sense that for the duration of the saccade (which is computed from the image distance it must move the fovea and the maximum velocity it can move), all other commands in the pipeline are overwritten, and the mode is changed to "saccade". Further commands trying to affect these instants may be ignored, depending on the (compile-time) policy desired.

NON-ZERO DELAY CONTROL

Slight amounts of delay destabilized the simulated system, as expected (see the Experiments section below). Control with delays can be stabilized by turning down gains and slowing the response of the system, but its performance then suffers. Successful control with delays incorporates some form of prediction [Mar79]. The controller implemented in the simulation is a version of a Smith predictor [Smi57,Smi58], which is the basic idea behind most modern methods.

Smith's Principle is that the desired output from a controlled system with delay p is the same as that desired from the delay-free system, only delayed by the delay p. Let the delay be z^{-p}, the delay-free series controller be $C(z)$, the desired delay controller be $\tilde{C}(z)$ and the plant be $A(z)$. The delay-free system transfer function will be

$$\frac{CA}{1+CA}.$$

The delay system with its desired controller has transfer function

$$\frac{\tilde{C}Az^{-r}}{1+\tilde{C}Az^{-r}}.$$

But Smith's Principle is

$$\frac{\tilde{C}Az^{-r}}{1+\tilde{C}Az^{-r}} = \frac{CAz^{-r}}{1+CA}$$

This quickly leads to the specification for the controller $\tilde{C}$ in terms of C, A, and z^{-r}:

$$\tilde{C} = \frac{C}{1+CA(1-z^{-r})}.$$

This simple principle has spawned a number of related controllers, often arising from each other by simple block-diagram manipulation. Figure 1 is one block diagram of a Smith prediction controller, and it describes the implemented system in the simulator.

If the maximum delay of a controller in the system is T. The plant model is a pipeline of enough future robot states to reach time T into the future, updated and extended once a tick. Ideally the robot's state is predictable, since only the control commands act on it. Practically there may be some plant noise. In the work so far, the world prediction is simplified by assuming the world is static and that the robot does all the moving (navigation in a static environment). As part of the experiments, target motion was added to test the system's response to a false target model.

EXPERIMENTS

DELAY-FREE CONTROL

In all the simulations, the goal of the system is to put one or both of its eyes squarely on the target (at retinal position (0,0)) and keep them there. The head is always in an upright position, so pans rotate the cameras about a vertical world axis, tilts rotate the cameras about a horizontal axis. With a static head, pans induce image x motion upon a static, foveated target and tilts induce image y motion. In all the graphs of this section, the horizontal axis is time, and the vertical axis is pan and tilt error, or equivalently the image x and y position of the target. Each graph shows both left and right eye x and y errors, but often the y errors are superimposed since the tilt platform is common to both cameras. In every case there is "peripheral blur", which is modelled by adding, outside a small "fovea", uniform noise to the target (x, y) location, with standard deviation proportional to $1/d$, where d is the euclidean distance of (x, y) from the (0,0) point. The simulation does not use realistic time-constants and speeds, which instead are scaled so that interesting effects happen within a few ticks.

Figs. 2 and 3 illustrate the cumulative effect of simply superimposing control capabilities: each operates independently and their outputs are simply summed at the effectors. Delays are zero, latencies (except for saccades) unity. In these two figures tracking is by position error signal.

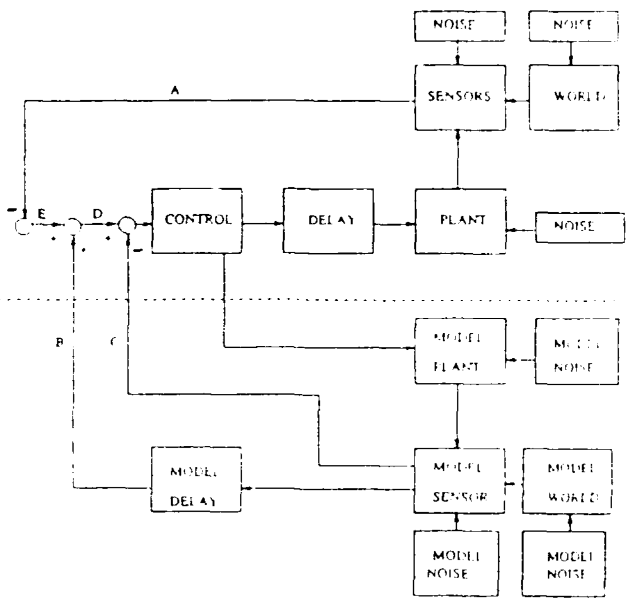

Figure 1: The implemented Smith predictor control. The block diagram is easily derived from the Smith predictor equation, with the MODEL PLANT, MODEL WORLD, and MODEL SENSOR blocks corresponding to A. $\tilde{C}$ is represented by the block labelled CONTROL and everything below the dashed line. The CONTROL block represents all five control systems, and the DELAY block represents a vector of their five independent delays. The PLANT, WORLD, and SENSOR blocks represent the robot simulation. Delayed control is implemented with a pipeline of controls to take place in the future, and the plant model is a similar pipeline of predicted robot states derived from the control.

(a) (b)

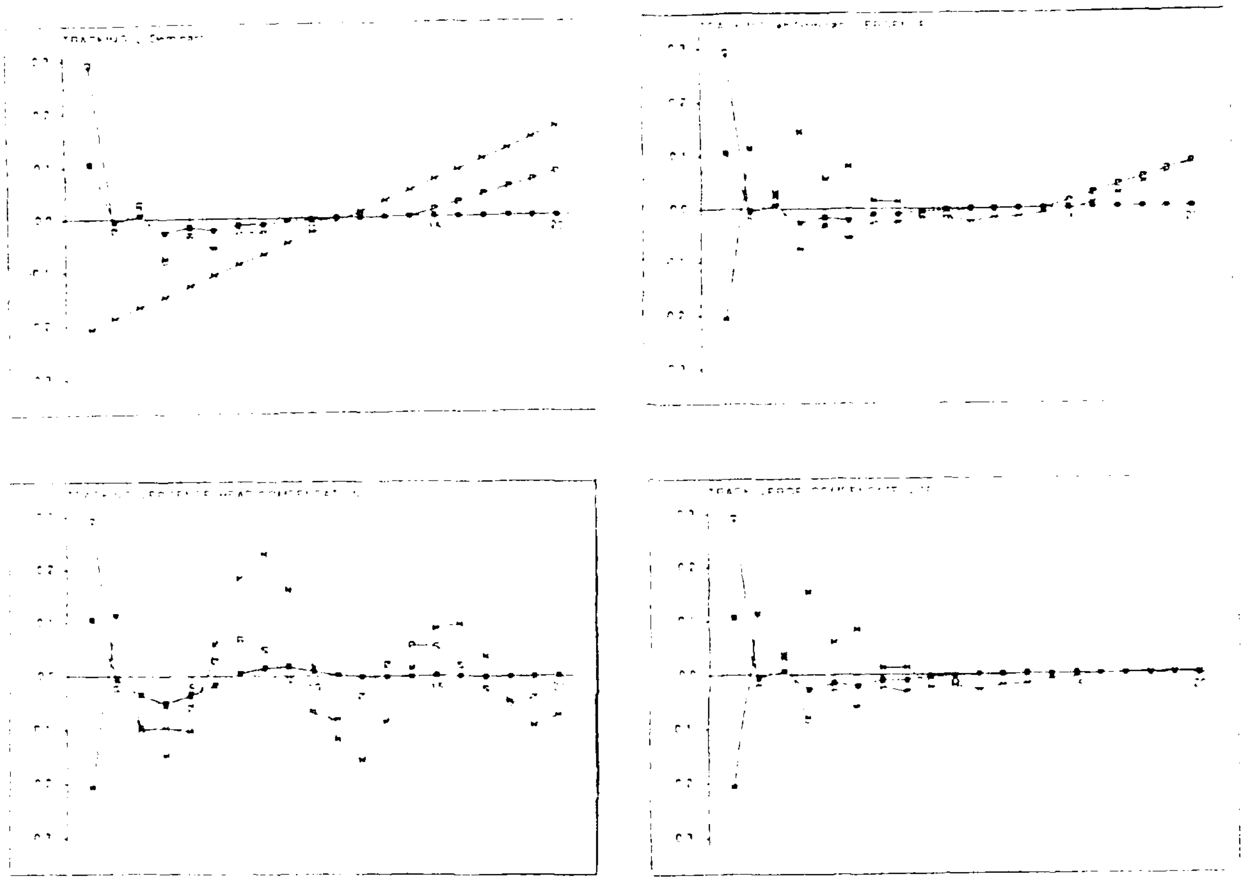

(c) (d)

Figure 2: Increasingly effective delay-free control results from superposition of noninteracting controllers. (a) Tracking only: The left (dominant) eye pans and tilts, inducing tilt in the right eye. The tracker uses a position error signal. The right eye gets no pan signal, and its horizontal error accrues from target motion. The left eye tracks successfully until it hits mechanical stop at tick 14. (b) Add vergence: Both eyes hit stops at about tick 15. (c) Add head compensation: This control is to keep eyes from hitting mechanical stops by turning the head in the same direction as the tracking motion. A less-desirable effect is to amplify the tracking signal, overcompensating and destabilizing the tracking. (d) Add VOR, which effectively compensates the head rotation with eye rotations.

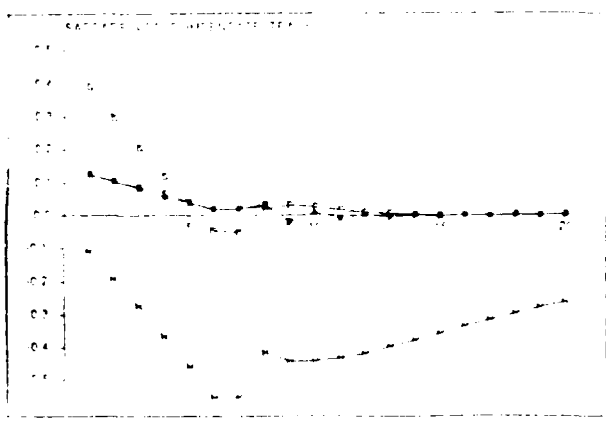

(a)

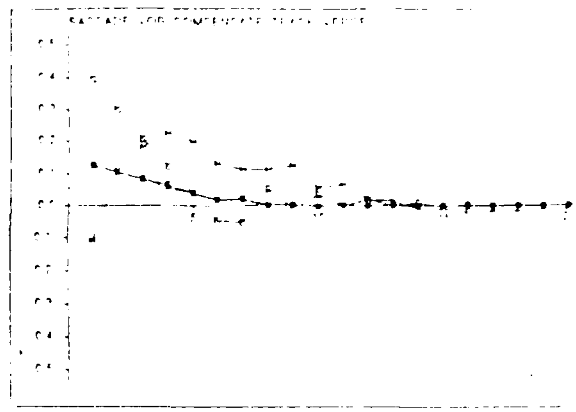

(b)

Figure 3: (a) Continuing the previous figure with tracking driven by position error, add saccades in which vergence, VOR, and head compensation are turned off during saccade. The saccade drives the left eye error more or less to zero (it is affected by the peripheral blurring effect which makes the initial location of the target image uncertain). It slews the right eye off target. When VOR, head compensation and vergence are turned on after the saccade, the first two reflexes have a transient effect. (b) Here let vergence run during the saccade but inhibit VOR and head compensation until after saccade completes.

Fig. 4 shows the effects of tracking with a velocity error signal. Here saccades are initiated if the target falls outside a fixed distance (here .1) from the fovea.

Finally, Fig. 5 shows the effects of control delay on the system. The smallest delays, applied uniformly or to just one control, destabilize the system seriously.

DELAY CONTROLS

As derived, the Smith predictor is appropriate for a single system control (or sensing) delay. In our system there will be differing delays reflecting different software actions (serial line plus VAL-II software versus VME bus connection to the eye motor controllers, for instance). The idea of the Smith predictor is easily extended, however.

Independent Delay Control

Two types of control were implemented using the Smith controller of Fig. 1. In the first, the controllers are ignorant of the delays of other controllers, and also ignorant of the sharing of output variables between controllers. Each controller knows its own delay T, and uses the following algorithm. *Look ahead time T and retrieve the predicted robot and control states for that time. Apply the control appropriate for these future states now.*

Fig 6 shows some sample effects of this independent delay-control strategy. The system is stable for certain combinations of delays, but is unstable unless all the non-vergence delays are the same.

Interacting Delay Control and Noise

The independent delay control algorithm is not as smart as it could be. The short-delay controls do not look into the future as far as the long-delay controls, and therefore they do not anticipate the effects of slower controls. This effect shows up when long-delay and short-delay controls affect each other's output, either directly or through the kinematic chain. The reason the verge reflex can run with different delay and not destabilize the independent delay control system is that no other control (barring saccade) affects the right camera's pan velocity, and panning is at the end of the kinematic chain. Assume each controller knows its

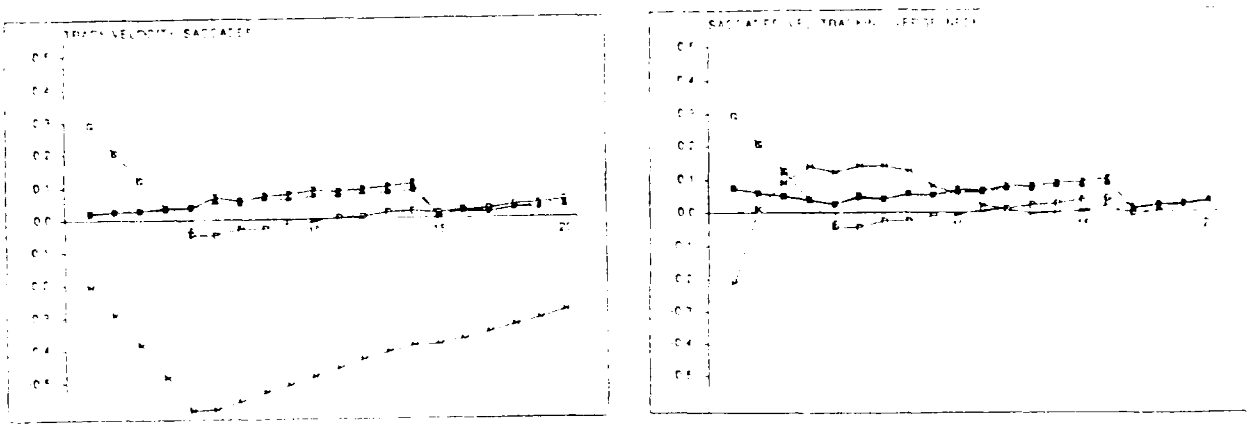

(a) (b)

Figure 4: (a) No vergence, velocity error tracking with saccades for position control. Tracking is subject to steady state position error. (b) Add vergence, and also change head kinematics (unknown to any controllers) from a "spherical" geometry to the Rochester robot's configuration of pan, tilt, and optic axes. The changed geometry has little effect.

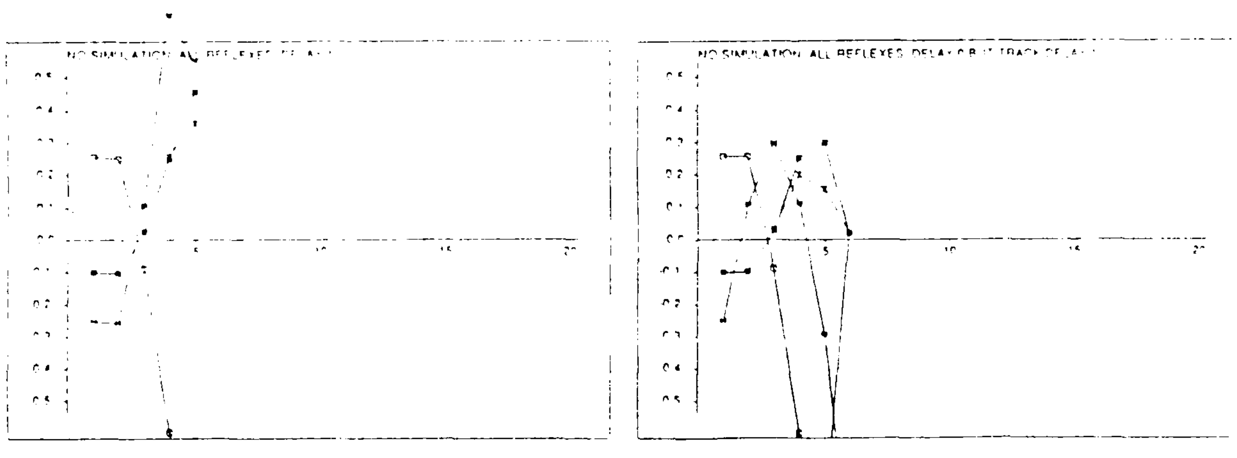

(a) (b)

Figure 5: (a) The no-delay controller applied to the system with a constant delay of one tick in all controls. Ideally this graph should be a delayed version of Fig. 2(d). (b) The no-delay controller applied with zero delay in all controls except tracking, which has a delay of one tick.

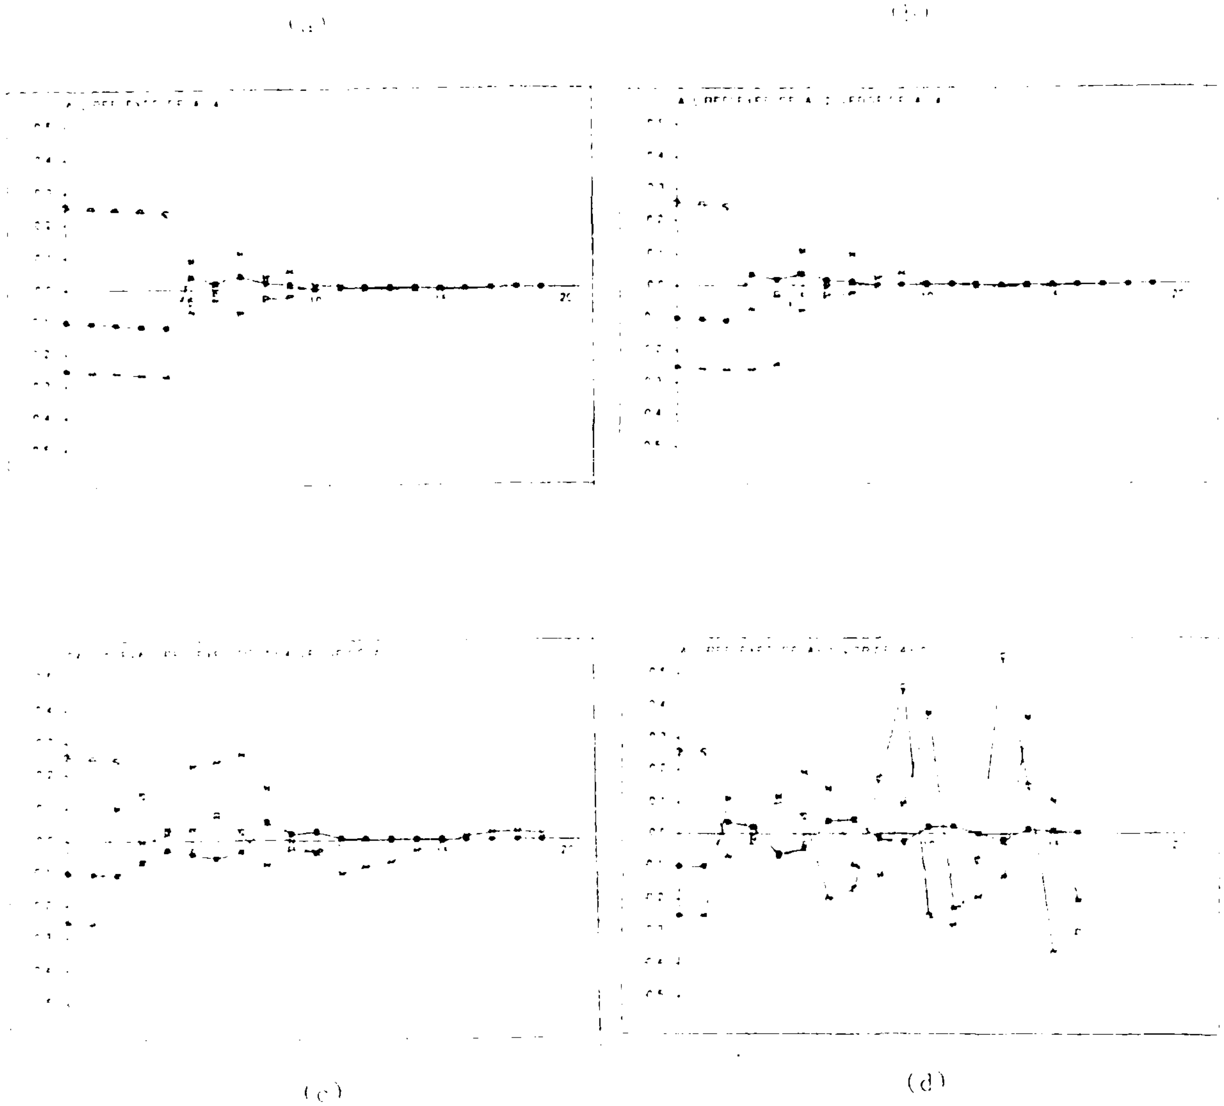

Figure 6. (a) To be compared with Fig. 2(d) and Fig. 5(a). The Smith predictor with independent control is stable with uniform controller delays. (b) Independent control also is stable with vergence control delay different. (c) Saccades induce transients but the system is still stable even if vergence delay different. (d) System is unstable if a non-vergence control, here VOR, has different delay from other non-vergence controls.

own delay T, and the delays of all the other controllers in the set $\{S\}$ that share an output with it. Then each controller can use the following (interacting controls) algorithm. *Look ahead the maximum delay M of any controller in $\{S\}$ and retrieve the predicted robot and control states for that time. Apply the control appropriate for these future states at (possibly future) time M-T.* This algorithm successfully copes with a different delay for each control (Fig. 7a).

An easy implementation of this algorithm that loses some flexibility is simply to increase the delay of all controls that share an output to be the maximum delay of any of their number and apply the independent delay control algorithm. Then all controls in the set look ahead as far as their slowest member, and act at the current moment. The resultant slowing of fast controls is of course suboptimal when they do not have to act in concert with slow controls.

Figures 7 and 8 show some experiments with interacting delay control, and introduce stochastic disturbances in the inputs and delays. The system is robust against sensor noise, or varying uncertainty in target location. The preliminary conclusion is that the system destabilizes with unpredictable delays when the outputs are changing relatively fast, but (of course) is less susceptible to unpredictable delays if the control outputs are only changing slowly.

DISCUSSION AND FUTURE WORK

SIMULATION AND REALITY

The goals for the simulator were to provide a kinematic and imaging model fairly close to that of the Rochester robot. The model has no dynamics, but neither does the robot from the point of view of the applications programmer; the current robot and motor control software hides this level. The simulator does seem adequate to illustrate the characteristics of different styles of control and to demonstrate the qualitative behavior resulting from control interaction, delays, and various forms of uncertainty. As the sophistication of the control technology at Rochester increases, a useful simulator would have to incorporate increasingly sophisticated models.

Likewise the simulator's exterior world and image-processing model is simple, consisting of a single point whose image is instantaneously and reliably (if noisily) found. To some extent this is also realistic, since it reflects the capability of frame-rate feature detection [Bro88], but it ignores the existence of more sophisticated operations or those with longer time-constants.

Simulation is likely to remain a basic tool in a real-time robotics laboratory, but as the control and visual environment gets sophisticated the simulations become slow and costly. The advent of cheap real-time hardware makes it increasingly practical to replace simulations with real-world experiments, which are more likely to yield relevant results.

COMPARISON WITH PRIMATE GAZE CONTROL MODELS

Because of its experimental accessibility, the simplicity of the plant involved, and the diverse collateral knowledge about the visual system, the gaze control system is the best-studied biological sensorimotor control system. The animal model most relevant to our robotic work is the primate, because of the close relationship of visual attention with fixation that arises with foveal (i.e. narrow-angle, high-resolution) vision. Gaze control in the cat and rabbit (and frog) is significantly different.

Knowledge of the primate gaze-control system might help provide insight to robot designers, and if the right hardware were available robotic equipment might be used to implement computational models of gaze control, thus providing an experimental facility complementary to the usual psychophysical and neuroscientific ones. The work described here is not yet dedicated to modeling biological systems, but nonetheless comparisions

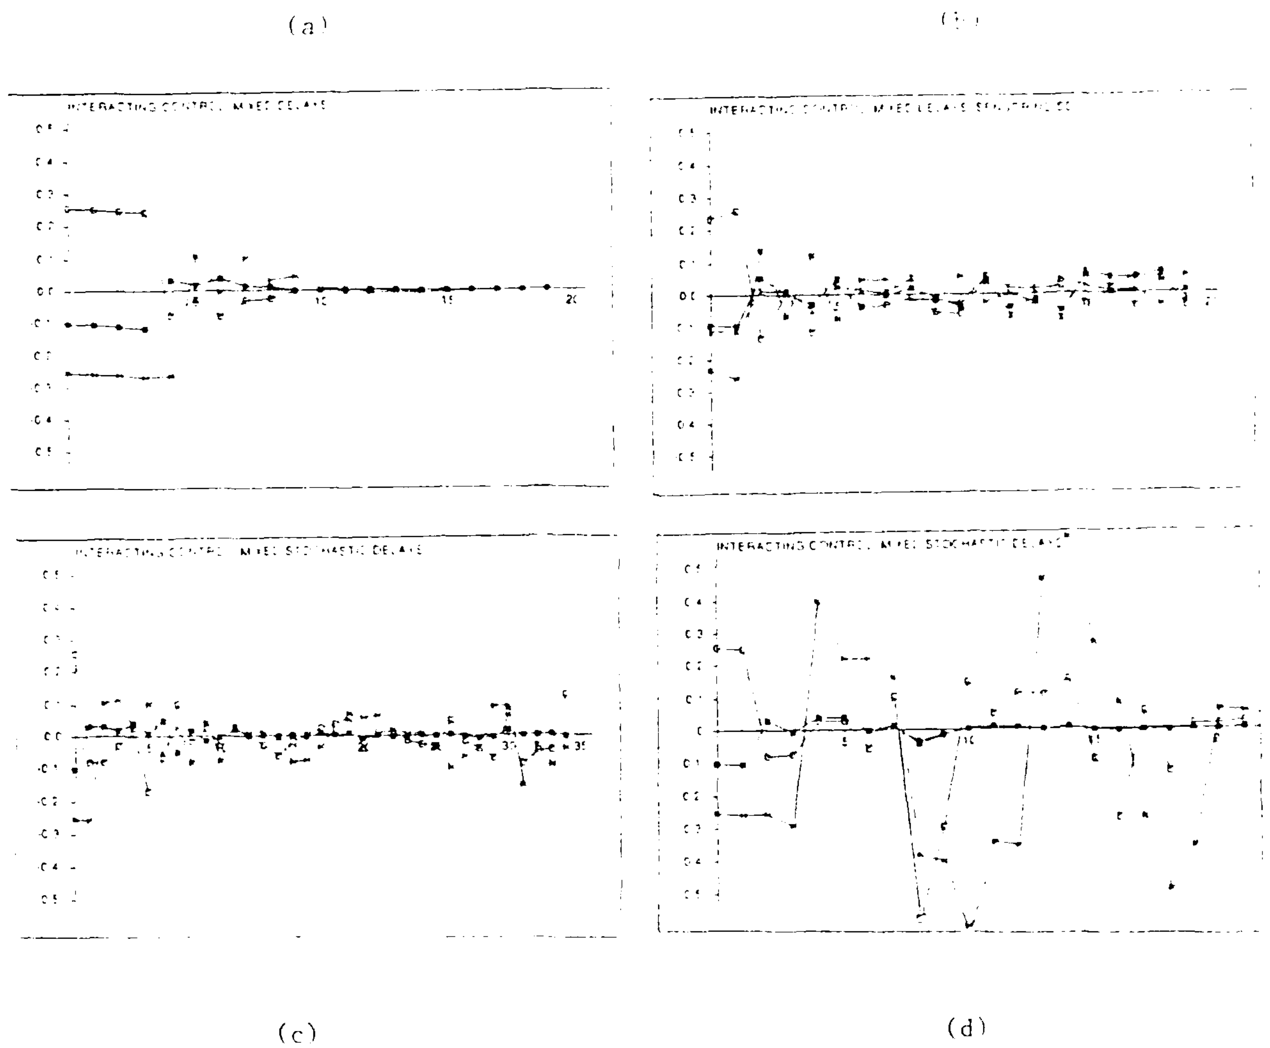

Figure 7: (a) The interacting control algorithm dealing successfully with a mixed set of delays. Here the longest non-vergence delay is three ticks, and the resultant behavior is that of a system whose non-vergence controls have a uniform delay of that amount. (b) Sensor noise (uniformly distributed disturbance of the target (x,y) location in each eye with $\sigma = 0.02$ in each dimension) does not affect stability, but causes excursions larger than its σ through the interaction of tracking and verging. (c) Here with probability .1 a control signal is delivered one tick early, and with probability .25 it is delivered one tick late. The system is on the verge of instability. (d) With same probabilities as in (c), more disturbances happen to occur early in the sequence when outputs are changing rapidly, destabilizing the system.

(a) (b)

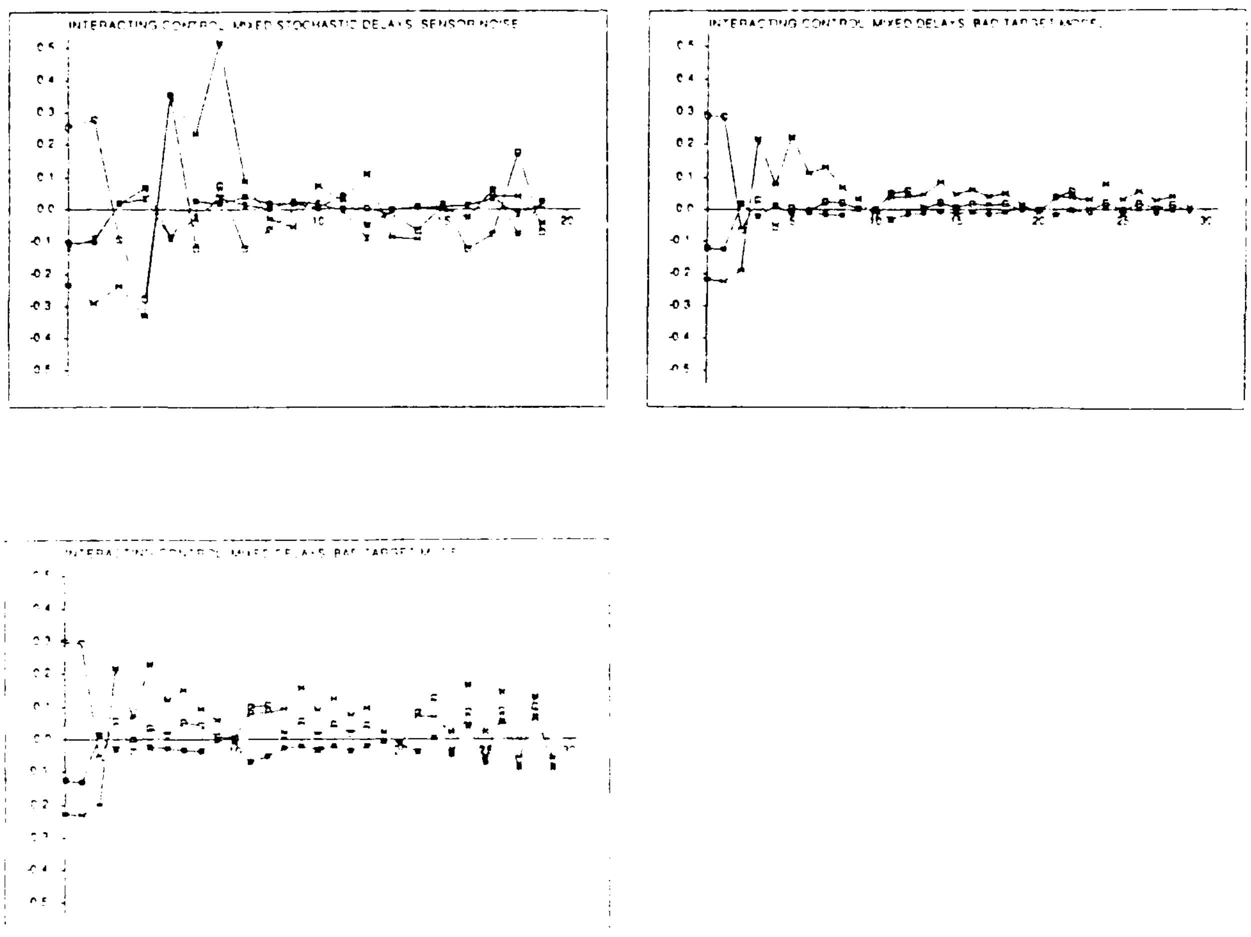

(c)

Figure 8: (a) Continuing from the previous figure, the previous sensor noise is added to the system along with the previous stochastic delays: the system is stable. (b) Here there is no noise (other than peripheral blurring), but the target model is wrong. The target is moving approximately perpendicular to the robot's motion instead of remaining static. The error periodicity of 10 ticks is interesting. (c) Here the situation is as in (b), but the target is moving faster, and toward the robot. As it gets close the controls cannot respond fast enough and the system destabilizes.

are inevitable, amusing, and possibly useful. This section is a very brief and admittedly selective sampling from the immense and rich (i.e. confusing and contradictory) literature on gaze and head control in biological systems. It seems fair to say that most of these systems interact, and that it is very difficult to lay down hard and fast rules about what individual systems can and cannot achieve.

Pursuit and Opto-Kinetic Reflex

The Opto-Kinetic Reflex (OKR) causes the eyes to follow a motion of the full visual field, and is driven (to first order) by "retinal slip", or optic flow. In primates the OKR comes in two stages, a faster (direct) and a slower (indirect), with the direct being more dominant in man. The smooth pursuit mechanism is to track small targets, and is often described as being driven by foveal retinal slip. Thus these two facilities are similar, and there is some thought that the direct part of the OKR response is just the smooth pursuit system [Col85].

The situation with smooth pursuit is anything but simple, however. It seems to be possible to pursue extra-foveal targets smoothly. Smooth eye movements cannot normally be induced without a smoothly-moving stimulus, but they persist after a target disappears, thus arguing that some form of prediction can excite the response [Eck83]. Smooth pursuit gain drops with stimulus velocity. Last, smooth pursuit in monkeys seems to be driven (in a large fraction of individuals) not just by velocity error but also by position and acceleration errors. Thus a model such as Young's (see below) that suggests a reconstructed target velocity is the control input (rather than a sensed optical flow) could be augmented with a broader range of error signals [LMT85].

The simulator has implemented both velocity control and position control with predictable results (compare Fig. 3(b) with Fig. 4(b)). Without position feedback, the system matches velocity and relies on saccades, which take place when position error goes over a threshold, for position control. There seems no advantage to this implementation unless optic flow velocity can be sensed directly, as opposed to position. For instance, if motion blur could be directly sensed, it would make a direct optic-flow velocity signal. Of course analysis of a particular motion-blur track could yield its centroid or endpoints, bringing us back to position control.

Vergence and Saccades

The primate vergence system is rather slow, and coupled to the focussing (accommodative) systems and the saccadic system. Vergence and accomodation are coupled pairwise, and the "near triad" is a reflex made up of these three systems, in which focus and vergence are both driven in the proper direction and faster than normal when a saccade from close to distant target (or the reverse) is made [Mil85].

Work with the Rochester robot has concentrated on "gross vergence", mediated through disparity computed between full-field images with variants of the cepstral filter [OP89]. The simulator described here is driven by horizontal disparity between the left and right target images. In the simulator, (which does not include focus) the cooperation of vergence and saccades is achieved simply, by the device of letting imaging, disparity calculation, and vergence reflex run during saccades. This method may or may not be nonbiological (as usual there is some dispute about the amount of visual processing that goes on during saccades). Its practical disadvantage is that it is inefficient: It is just as easy to have the saccade control both eyes. The only reason the current simulator does not run this way is that it is less interesting.

The saccadic system has a longer delay than smooth pursuit (120ms as opposed to 50 ms), reflecting its higher-level control origins. It can move the eye at 300 to 400 degrees/second. It is often modeled as a sampled-data system, kept stable by a latency and trigger mechanism that inhibits its firing again before the system has settled. In our robot system, saccades should not be needed for position control during tracking, and thus will be associated with shifts of attention, or at least of visual resource commitment.

In the experiments shown, the maximum saccade speed was limited but the maximum speeds for other reflexes were not (compare the .1 rad/tick saccade rate in Fig. 3(a) with the .3 rad/tick speed of the tracking and vergence in Fig. 2(d). Clearly the control should not be allowed to command unrealistic

speeds, and the relative strengths of the outputs must be adjusted. In our simulation, the strictly "left eye dominant" implementation of saccades and of tracking is almost certainly an exaggeration of the ocular dominance effects in primates. Still, from a practical point of view it means that the necessary low-level vision computations do not need to be carried out in both eyes simultaneously.

The Vestibulo-Ocular Reflex

The Vestibulo-Ocular Reflex (VOR) stabilizes gaze by counteracting commanded head movements with eye movements. It is the fastest visual reflex, with a delay of only approximately 16 milliseconds. It is an open-loop control, in the sense that vestibular sensor output is converted to eye muscle input and delivered through a path of approximately three synapses. It can be a high gain control (gain approximately 1). it can often exactly cancel out head motion effects. The VOR being open loop, there is a general problem of how it internally models the system it is controlling.

Research on the VOR has addressed the geometrical aspect of its modelling: the conversion of sensor signals in the coordinate systems of the semicircular canals to effector signals for the variously-placed eye muscles. Robinson [Rob85] models the geometrical transformations as 3x3 matrices operating on 3-vectors. Changing matrix components can accomplish adaptation, and the adaptation can be driven by stimuli such as retinal slip (indicating a failure of the reflex) without explicitly modelling the sensorimotor system. Pellionisz [Pel85,PP88] uses tensors to model the differing transformation properties of the sensory and motor vectors and transformations, and addresses the problem of underdetermined control of the many muscles that accomplish eye and head movements by the relatively small number of sensor dimensions.

The VOR's input originates in the linear and angular accelerometers of the otolith organs and semicircular canals. They have very short time constants, but the VOR operates correctly for slow velocities. This leads to the postulation of a "velocity storage mechanism" that integrates the output of the accelerometers and makes the resulting velocity signal available for control (e.g. [RC85]).

Other VOR work addresses its time-dependent behavior: its gain and phase-lag characteristics under different conditions (e.g. several papers in [BJ85]). Much of the VOR's behavior can be explained as parameter variation among its gain, bias, and time constants. Miles *et al.* [MOL85] develop a multi-channel model to explain VOR's ability to cope with the frequency-dependent output characteristics of the sensors, with frequency-selective adaptation properties of the VOR itself, and with other adaptive properties of the VOR. This work presents explicit transfer functions for the semicircular canals, the oculomotor plant, the velocity storage mechanism, and the neural channels that convert head velocity estimates to motor outputs. The channel model is linear and can be stated as a lumped-parameter linear system, but the channels make it easier to identify which gains must be changed to reduce system errors.

A basic aspect of the VOR is its adaptability. The reflex adapts over time to changes in the optical system (e.g. artificially induced dysmetria) [Rob85]. The VOR interacts with other reflexes and the stimuli that evoke them. For example, large-field rotations that elicit the OKR have an interesting effect. If they are slow, they bias the VOR (and the opto-kinetic system) in the same direction, which tends to cancel the movement effect. If they are fast, they induce effects in the opposite direction, which may be interpreted as ignoring the movement effect [Col85]. VOR gain can be depressed from 1.0 to 0.1 by training that involves no visual input (subject imagines tracking a target attached to head while moving head in the dark), and is likewise significantly affected by verbal instructions and other seemingly unrelated activities (such as mental arithmetic) [JB85].

Adaptation and modeling can come together in VOR behavior that adapts to repetitive patterns (a perhaps familiar example is disembarking from a longish sailing journey). One way to achieve this capability is through a "pattern storage" mechanism that effectively produces and uses a model of the outside world. Some workers are attracted to this idea, others seem to think it is unnecessary and are explicable by, for instance, channel adaptation.

What has all this to do with a robotic VOR? Many of the issues mentioned above can be made to vanish

We may know the relation of the sensor output to the desired motor output if we decide to model the robot and head kinematics accurately. (In fact in the simulation, the robotic VOR makes several approximations including a "spherical" geometry for the camera rotation axes, a small-angle approximation, and others.) We can sense velocities directly or even actively monitor the relevant control signals we need to cancel. The fundamental issues that still need significant work involve adaptation and interaction. Adequate understanding of these issues would not only give the robot system the efficiency exhibited by natural systems, but could mean that such exercises as accurate kinematic modeling would become unnecessary.

Head Control

There is less written on head control than on gaze control, but a good recent collection of work exists [PR88]. There are various head stabilization reflexes, some tied to optical stimulation. The relation of head control strategies to the evolution of particular brain mechanisms and the existence of foveate vision is explored by Roucoux and Crommelinck [RC88]. Some fairly detailed biomechanical head models exist, and head movements have been investigated from the point of view of optimal control theory. Head movements can be quite rapid (600-700 degrees/second) and are part of normal long-distance saccades in primates. Thus the saccadic and head-control system work together to achieve gaze redirection. There has been some work here (e.g. [Gui88]) indicating that head movements can take place at differing times relative to saccades. Typically, they lead or lag depending on whether the target location is predictable or not.

This coupling of head and eye movements is clearly more sophisticated than the compensatory reflex implemented in the simulation, which is not coupled to saccades at all and which must lag eye movements since it is only driven by eye positions. Thus more work needs to be done if we are to achieve the increased rapidity of gaze redirection that arises when both head and eyes are moved in a coordinated way.

Another Model of Delay Control

The control scheme implemented in this simulation, the Smith predictor, differs from a scheme seemingly first proposed in a gaze-control context by Young, taken a step further by Robinson, and used recently in robotic gaze-control for an agile, two-eyed robotic head at Harvard University [CF88].

Young [You77] wanted to explain how smooth pursuit avoided instability in the presence of two difficulties that apply if tracking is modeled as a pure negative feedback system. First, the error, and thus control, signal is zero when accurate tracking is achieved; this should send eye velocity transiently to zero. Second, tracking performance is better than it should be given the delays in the control loop and the time constants of the processes. His proposal is that the system tracks not the retinal image, but a neural signal that corresponds to target motion (in the world).

In 1971 (for a recent reference, applied to saccadic, tracking, and limb control, see [Rob88]) Robinson proposed a mechanism to implement Young's idea. In the negative feedback system the eye velocity is fed back and subtracted from the target velocity (with some delay). If the eye is in the process of tracking, then the target velocity is the sum of the eye velocity (with respect to the head) and the target's retinal velocity (its velocity with respect to the eye). But the latter is just the error signal resulting from negative feedback. Thus an estimated target velocity signal can be constructed by positively feeding back the commanded eye motion into the control loop, delayed to arrive at the proper time to combine with the error term produced by negative feedback. This mechanism not only provides a signal based on the target's true motion, but it cancels the negative feedback and thus removes the possibility of oscillations.

Robinson's scheme is related to the Smith controller shown in Figure 1 in the following way. In Figure 1, the signal at E is an error signal, and the one at D is a difference of error signals that is zero when perfect tracking is taking place. This difference of errors is a delayed (but consistent) error signal that is added to the predicted error signal in the non-delayed path C. The controller in Figure 1 tries to drive errors to zero. To change Figure 1 to Robinson's scheme, delete path C and remove the modelled world and sensor from the lower half of the block diagram. Then path B carries the simulated plant, not the simulated error. Path E still contains error, but path D now contains a prediction, or reconstruction, of the world state. Thus

the controller now must treat the signal at D as a set point to be achieved through open-loop methods, not as an error. Robinson proposes parametric adaptive control (in the form of two related gains) to provide adaptative capability should the open loop yield the wrong results.

There are thus some similarities between the two schemes, but the underlying control philosophies are rather different. In paricular, losing the power of negative feedback is a large sacrifice that the roboticist may not need to make. The Smith predictor control system keeps the advantage of feedback control (running on the modelled world and plant). There are many methods of estimation, observation, and prediction of world, sensor, and plant used in modern control theory, and thus the Smith model allows for flexibility in the assumptions underlying its predictions.

FUTURE WORK

We plan to supply more quantitative model parameters, and to try to model the spatial and temporal scales that actually apply in the laboratory. Sensitivity analysis will be undertaken to quantify the effects of various disturbances, especially the probiem of unpredictable delays.

We plan to integrate some of the existing Kalman filtering tracking utilities [Bro89,BF88] to perform estimation of the target's state. Also we may explore estimation techniques [Gel73,Ber76,Eyk74] instead of simulation techniques to predict the state of the plant.

The simulated system can support other relevant aspects to the control problem, including the important one of adapting to changes in the plant. In other work, we have implemented "the MIT rule", which is a gradient descent method similar to back-propagation learning in neural nets, to learn part of the robot head geometry. In a way this learning system acts like another control system, with inputs the discrepencies between expected and observed target motions given eye motions, and outputs are parameters to the modeled plant (in this case, lengths of links in the head kinematic chain).

Implementation of an increasingly sophisticated gaze control system on the Rochester robot should take place over the next few years. We anticipate substituting a Butterfly Parallel Processor with multiple input and output ports for the central controller of the system.

Acknowledgements

This work is funded by the DARPA U.S. Army Engineering Topographic Laboraties Grant DACA76-85-C-0001 and the Air Force Systems Command (RADC, Griffiss AFB, NY) and Air Force OSR Contract F30602-85-C-0008, which supports the Northeast Artificial Intelligence Consortium. The assistance of Prof. J. Michael Brady, Dr. Hugh Durrant-Whyte, Dr Ron Daniels, and the RRG research and administrative staff is gratefully acknowledged.

References

[Ber76] Dimitri I. Bertsekas. *Dynamic Programming and Stochastic Control*. Academic Press, 1976.

[BF88] Y. Bar-Shalom and T.E. Fortmann. *Tracking and Data Association*. Academic Press, 1988.

[BJ85] A. Berthoz and G. Melvill Jones. *Adaptive Mechanisms in Gaze Control: Facts and Theories*. Elsevier, 1985.

[BO88] D. H. Ballard and A. Ozcandarli. Real-time kinetic depth. In *Second Int. Conf. on Computer Vision*, November 1988.

[BR88] C. M. Brown and R. Rimey. *Kinematics, Coordinate Systems and Conversions for The Rochester Robot*. Technical Report 255, University of Rochester, September 1988.

[Bro88] C. M. Brown. *The Rochester Robot*. Technical Report 257, University of Rochester, September 1988.

[Bro89] C. M. Brown. Kalman filtering for tracking and control. In *DARPA Image Understanding Workshop*, June 1989.

[CF88] J. J. Clark and N. J. Ferrier. Modal control of an attentive vision system. In *Second Int. Joint Conference on Computer Vision*, November 1988.

[Col85] H. Collewijn. Integration of adaptive changes of the optokinetic reflex, pursuit and the vestibulo-ocular reflex. In A. Berthoz and G. Melvill Jones, editors, *Adaptive Mechanisms in Gaze Control*. Elsevier, 1985.

[Eck83] R. Eckmiller. Neural control of foveal pursuit versus saccadic eye movements in primates – single-unit data and models. *IEEE Trans. on Syst., Man, and Cyber.*, SMC-13(5):980 – 989, Sept./Oct. 1983.

[Eyk74] Pieter Eykhoff. *System Identification: Parameter and State Estimation*. Wiley and Sons, 1974.

[Gel73] Arthur C. Gelb. *Applied Optimal Estimation*. The MIT Press, 1973.

[Gui88] D. Guitton. Eye-head coordination in gaze control. In B. W. Peterson and F. J. Richmond, editors, *Control of Head Movement*, Oxford University Press, 1988.

[JB85] G. Melvill Jones and A. Berthoz. Mental control of the adaptive process. In A. Berthoz and G. Melvill Jones, editors, *Adaptive Mechanisms in Gaze Control*, Elsevier, 1985.

[LMT85] S. G. Lisberger, E. J. Morris, and L. Tychsen. Visual motion processing and sensory-motor integration for smooth pursuit eye movements. In A. Berthoz and G. Melvill Jones, editors, *Adaptive Mechanisms in Gaze Control*, Elsevier, 1985.

[Mar79] J. E. Marshall. *Control of Time-Delay Systems*. Peter Peregrinus Ltd., 1979.

[Mil85] F. A. Miles. Adaptive regulation in the vergence and accommodation control systems. In A. Berthoz and G. Melvill Jones, editors, *Adaptive Mechanisms in Gaze Control*, Elsevier, 1985.

[MOL85] F. A. Miles, L.M. Optican, and S. G. Lisberger. An adaptive equalizer model of the primate vestibulo-ocular reflex. In A. Berthoz and G. Melvill Jones, editors, *Adaptive Mechanisms in Gaze Control*. Elsevier, 1985.

[OP89] T. Olson and R. Potter. Real-time vergence control. In *Computer Vision and Pattern Recognition 1989*, June 1989.

[Pel85] A. J. Pellionisz. Tensorial aspects of the multidimensional approach to the vestibulo-oculomotor reflex and gaze. In A. Berthoz and G. Melvill Jones, editors, *Adaptive Mechanisms in Gaze Control*, Elsevier, 1985.

[PP88] A. J. Pellionisz and B. W. Peterson. A tensorial model of neck motor activation. In B. W. Peterson and F. J. Richmond, editors, *Control of Head Movement*, Oxford University Press, 1988.

[PR88] B. W. Peterson and F. J. Richmond. *Control of Head Movement*. Oxford University Press, 1988.

[RC85] T. Raphan and B. Cohen. Velocity storage and the ocular response to multidimensional vestibular stimuli. In A. Berthoz and G. Melvill Jones, editors, *Adaptive Mechanisms in Gaze Control*, Elsevier, 1985.

[RC88] A. Roucoux and M. Crommelinck. Control of head movement during visual orientation. In B. W. Peterson and F. J. Richmond, editors, *Control of Head Movement*, Oxford University Press, 1988.

[Rob85] D. A. Robinson. The coordinates of neurons in the vestibulo-ocular reflex. In A. Berthoz and G. Melvill Jones, editors, *Adaptive Mechanisms in Gaze Control*, Elsevier, 1985.

[Rob88] D. A. Robinson. Why visuomotor systems don't like negative feedback and how they avoid it. In M. A. Arbib and A. R. Hanson, editors, *Vision, Brain, and Cooperative Computation*, MIT Press, 1988.

[Smi57] O. J. M. Smith. Closer control of loops with dead time. *Chemical Engg. Prog. Trans.*, 53(5):217-219, 1957.

[Smi58] O. J. M. Smith. *Feedback Control Systems*. McGraw-Hill, 1958.

[You77] L.R. Young. Pursuit eye movement – what is being pursued? *Dev. Neurosci.: Control of Gaze by Brain Stem Neurons*, 1:29-36, 1977.